

Understanding Set Cover: Sub-exponential Time Approximations and Lift-and-Project Methods

Eden Chlamtác*
Ben Gurion University

Zachary Friggstad†
University of Waterloo

Konstantinos Georgiou†
University of Waterloo

January 28, 2013

Abstract

Recently, Cygan, Kowalik, and Wykurz [IPL 2009] gave sub-exponential-time approximation algorithms for the SET COVER problem with approximation ratios better than $\ln n$. In light of this result, it is natural to ask whether such improvements can be achieved using lift-and-project methods. We present a simpler combinatorial algorithm which has nearly the same time-approximation tradeoff as the algorithm of Cygan et al., and which lends itself naturally to a lift-and-project based approach.

At a high level, our approach is similar to the recent work of Karlin, Mathieu, and Nguyen [IPCO 2011], who examined a known PTAS for KNAPSACK (similar to our combinatorial SET COVER algorithm) and its connection to hierarchies of LP and SDP relaxations for KNAPSACK. For SET COVER, we show that, indeed, using the trick of “lifting the objective function”, we can match the performance of our combinatorial algorithm using the LP hierarchy of Lovász and Schrijver. We also show that this trick is essential: even in the stronger LP hierarchy of Sherali and Adams, the integrality gap remains at least $(1 - \varepsilon) \ln n$ at level $\Omega(n)$ (when the objective function is not lifted).

As shown by Aleknovich, Arora, and Turlakis [STOC 2005], SET COVER relaxations stemming from SDP hierarchies (specifically, LS_+) have similarly large integrality gaps. This stands in contrast to KNAPSACK, where Karlin et al. showed that the (much stronger) Lasserre SDP hierarchy reduces the integrality gap to $(1 + \varepsilon)$ at level $O(1)$. For completeness, we show that LS_+ also reduces the integrality gap for KNAPSACK to $(1 + \varepsilon)$. This result may be of independent interest, as our LS_+ -based rounding and analysis are rather different from those of Karlin et al., and to the best of our knowledge this is the first explicit demonstration of such a reduction in the integrality gap of LS_+ relaxations after few rounds.

*Research supported in part by an ERC Advanced grant. Email: chlamtac@cs.bgu.ac.il

†Email: {zfriggstad,k2georgiou}@math.uwaterloo.ca

1 Introduction

The SET COVER problem is one of the most fundamental and well-studied problems in approximation algorithms, and was one of Karp’s original 21 NP-complete problems [20]. It can be stated quite plainly: given a finite set X of n items and a collection $\mathcal{S} \subseteq 2^X$ of m subsets of X called “cover-sets”, the SET COVER problem on instance (X, \mathcal{S}) is the problem of finding the smallest collection \mathcal{C} of cover-sets in \mathcal{S} such that $X = \bigcup_{S \in \mathcal{C}} S$. That is, every item in X must appear in at least one cover-set in \mathcal{C} . Here, we will consider the minimum cost (or weighted) version of the problem, where each cover-set S has a nonnegative cost $c(S)$, and the goal is to find a collection of cover-sets with minimum total cost, subject to the above constraint.

As is well known, the problem can be approximated within a logarithmic factor. For instance, Johnson [18] showed that for uniform costs, the simple greedy algorithm that iteratively chooses the cover-set containing the maximum number of uncovered elements gives an H_n -approximation (where $H_n = \ln n + O(1)$ is the n ’th harmonic number $\sum_{k=1}^n 1/k$). Later, Lovász [22] showed that the cost of the solution found by this greedy algorithm is at most an H_n -factor larger than the optimum value of the natural linear programming (LP) relaxation.

Chvátal [10] extended these results to a greedy algorithm for the weighted case. He also proved that the approximation guarantee of this algorithm is actually only H_b where b is the size of the largest cover-set in \mathcal{S} , and moreover that this algorithm gives an H_b -factor approximation relative to the optimum of the natural LP relaxation (thus extending the integrality gap bound of Lovász [22]). Slavík [30] refined the lower-order terms in the analysis of the greedy algorithm’s approximation guarantee, giving a tight bound of $\ln n - \ln \ln n + O(1)$. Srinivasan [31] also improved the lower-order terms in the upper bound on the integrality gap. On the other hand, the integrality gap of the standard LP relaxation for SET COVER is at least $(1 - o(1)) \ln n$. Recently, Cygan, Kowalik, and Wykurz [11] demonstrated that Set Cover can be approximated within $(1 - \epsilon) \cdot \ln n + O(1)$ in time $2^{n^\epsilon + O(\log m)}$. It is interesting to note that this time-approximation tradeoff is essentially optimal assuming Moshkovitz’s Projection Games Conjecture [24] and the Exponential Time Hypothesis (ETH) [17].

From the hardness perspective, Feige [12] showed that for every constant $\epsilon > 0$, there is no $(1 - \epsilon) \ln n$ -approximation algorithm for SET COVER unless all problems in NP can be solved deterministically in time $n^{O(\log \log n)}$. To date, the strongest hardness of approximation for SET COVER assuming only $P \neq NP$ gives a $c \ln n$ -hardness for $c \approx 0.2267$ (Alon et al. [2]).

1.1 Hierarchies of convex relaxations and connection to Knapsack

One of the most powerful and ubiquitous tools in approximation algorithms has been the use of mathematical programming relaxations, such as linear programming (LP) and semidefinite programming (SDP). The common approach is as follows: solve a convex (LP or SDP) relaxation for the 0-1 program, and “round” the relaxed solution to give a (possibly suboptimal) feasible 0-1 solution. Since the approximation ratio is usually analyzed by comparing the value of the relaxed solution to the value of the output (note that the 0-1 optimum is always sandwiched between these two), a natural obstacle is the worst case ratio between the relaxed optimum and the 0-1 optimum, known as the *integrality gap*.

While for many problems, this approach gives optimal approximations (e.g., Raghavendra [26] shows this for all CSPs, assuming the UNIQUE GAMES Conjecture), there are still many cases where natural LP and SDP relaxations have large integrality gaps. This limitation can be circumvented

by considering more powerful relaxations. In particular, Sherali and Adams [28] Lovász and Schrijver [23], and Lasserre [21] each have devised different systems, collectively known as *hierarchies* or *lift-and-project techniques*, by which a simple relaxation can be strengthened until the polytope (or the convex body) it defines converges to the convex hull of the feasible 0-1 solutions. It is known that, for each of these hierarchies, if the original relaxation has n variables, then the relaxation at level t of the hierarchy can be solved optimally in time $n^{O(t)}$. Thus, to achieve improved approximations for a problem in polynomial (resp. sub-exponential time), we would like to know if we can beat the integrality gap of the natural relaxation by using a relaxation at level $O(1)$ (resp. $o(n/\log n)$) of some hierarchy.

Initially, this approach was refuted (for specific problems) by a long series of results showing that integrality gaps do not decrease at sufficiently low levels (see, e.g. [4, 1, 14, 6]). Positive results have also recently emerged (e.g. [8, 29, 5, 15]), where improved approximations were given using constant-level relaxations from various hierarchies. For a survey on both positive and negative results, see [9].

Recently, Karlin et al. [19] considered how the use of LP and SDP hierarchies affects the integrality gap of relaxations for KNAPSACK. This is of particular relevance to us, since their approach relies on a well-known PTAS for KNAPSACK which is similar in structure to our own combinatorial algorithm for SET COVER. They showed that, while the Sherali-Adams LP hierarchy requires $\Omega(n)$ levels to bring the integrality gap below $2 - o(1)$, level k of the Lasserre SDP hierarchy brings the integrality gap down to $1 + O(1/k)$. While we would like to emulate the success of their SDP-hierarchy-based approach (and give an alternative sub-exponential algorithm for SET COVER), we note that for SET COVER Alekhovich et al. [1] have shown that the SDP hierarchy LS_+ , due to Lovász and Schrijver, requires $\Omega(n)$ levels to bring the integrality gap below $(1 - o(1)) \ln n$. Nevertheless, the comparison is not perfect, since the Lasserre hierarchy is much stronger than LS_+ . In particular, previously it was not known whether LS_+ also reduces the integrality gap for KNAPSACK (and indeed, the algorithm of Karlin et al. [19] relied a powerful decomposition theorem for the Lasserre hierarchy which does not seem to be applicable to LS_+).

1.2 Our Results

To facilitate our lift-and-project based approach, we start by giving in Section 3.1 a simple new sub-exponential time combinatorial algorithm for SET COVER which nearly matches the time-approximation tradeoff guarantee in [11].

Theorem 1.1. *For any (not necessarily constant) $1 \leq d \leq n$, there is an $H_{n/d}$ -approximation algorithm for SET COVER running in time $\text{poly}(n, m) \cdot m^{O(d)}$.*

While this theorem is slightly weaker than the previous best known guarantee, our algorithm is remarkably simple, and will be instrumental in designing a similar lift-and-project based SET COVER approximation. The algorithm is combinatorial and does not rely on linear programming techniques. By choosing $d = n^\epsilon$, we get a sub-exponential time algorithm whose approximation guarantee is better than $\ln n$ by a constant factor.

Next in Section 3.2, we show that using level d of the linear programming hierarchy of Lovász-Schrijver [23], we can match the performance of the algorithm we use to prove Theorem 1.1, though only if the “lifting” is done after guessing the value of the objective function (using a binary search), and adding this as a constraint a priori. In this case, the rounding algorithm is quite fast, and avoids

the extensive combinatorial guessing of our first algorithm, while the running time is dominated by the time it takes to solve the LP relaxation.

On the other hand, without the trick of “lifting the objective function”, we show in Section 4 that even the stronger LP hierarchy of Sherali-Adams [28] has an integrality gap of at least $(1-\varepsilon) \ln n$ at level $\Omega(n)$. Specifically, we show the following

Theorem 1.2. *For every $0 < \varepsilon, \gamma \leq \frac{1}{2}$, and for sufficiently large values of n , there are instances of SET COVER on n cover-sets (over a universe of n items) for which the integrality gap of the level- $\lfloor \frac{\gamma(\varepsilon-\varepsilon^2)}{1+\gamma} n \rfloor$ Sherali-Adams LP relaxation is at least $\frac{1-\varepsilon}{1+\gamma} \ln n$.*

As we have mentioned, the prospect of showing a positive result using SDP hierarchies is unlikely due to the work of Alekhovich et al. [1] which gives a similar integrality gap for LS_+ .

For completeness, we also show in Section 5 that the Lasserre-hierarchy-based KNAPSACK algorithm of Karlin et al. [19] can be matched using the weaker LS_+ hierarchy and a more complex rounding algorithm. Specifically, we show that the integrality gap of the natural relaxation for KNAPSACK can be reduced to $1 + \varepsilon$ using $O(\varepsilon^{-3})$ rounds of LS_+ . This highlights a fundamental difference between KNAPSACK and SET COVER, despite the similarity between our combinatorial algorithm and the PTAS for KNAPSACK on which Karlin et al. rely. Formally, we prove the following:

Theorem 1.3. *The integrality gap of level k of the LS_+ relaxation for KNAPSACK is at most $1 + O(k^{-1/3})$.*

In what follows, and before we start the exposition of our results, we present in Section 2 the Lovász-Schrijver system along with some well-known facts that we will need later on. We end with a discussion of future directions in Section 6.

2 Preliminaries on the Lovász-Schrijver System

For any polytope P , this system begins by introducing a nonnegative auxiliary variable x_0 , so that in every constraint of P , constants are multiplied by x_0 . This yields the cone $K_0(P) := \{(x_0, x_0 \mathbf{x}) \mid x_0 \geq 0 \text{ \& \; } \mathbf{x} \in P\}$. For an n -dimensional polytope P , the Lovász-Schrijver system finds a hierarchy of nested cones $K_0(P) \supseteq K_1(P) \supseteq \dots \supseteq K_n(P)$ (in the SDP variant, we will write $K_t^+(P)$), defined recursively, and which enjoy remarkable algorithmic properties. In what follows, let \mathcal{P}_k denote the space of vectors indexed by subsets of $[n]$ of size at most k , and for any $\mathbf{y} \in \mathcal{P}_k$, define the moment matrix $Y^{[\mathbf{y}]}$ to be the square matrix with rows and columns indexed by sets of size at most $\lfloor k/2 \rfloor$, where the entry at row A and column B is $y_{A \cup B}$. Also we denote by $\mathbf{e}_0, \mathbf{e}_1, \dots, \mathbf{e}_n$ the standard orthonormal basis of dimension $n+1$, such that $Y^{[\mathbf{y}]} \mathbf{e}_i$ is the i -th column of the moment matrix.

Definition 2.1 (The Lovász-Schrijver (LS) and Lovász-Schrijver SDP (LS_+) systems). *Consider the conified polytope $K_0(P)$ defined earlier (let us also write $K_0^+(P) = K_0(P)$). The level- t Lovász-Schrijver cone (relaxation or tightening) $K_t(P)$ (resp. $K_t^+(P)$) of LS (resp. LS_+) is recursively defined as all $n+1$ dimensional vectors $(x_0, x_0 \mathbf{x})$ for which there exist $\mathbf{y} \in \mathcal{P}_2$ such that $Y^{[\mathbf{y}]} \mathbf{e}_i, Y^{[\mathbf{y}]} (\mathbf{e}_0 - \mathbf{e}_i) \in K_{t-1}(P)$ (resp. $K_{t-1}^+(P)$) and $(x_0, x_0 \mathbf{x}) = Y^{[\mathbf{y}]} \mathbf{e}_0$. The level- t Lovász-Schrijver SDP tightening of LS_+ asks in addition that $Y^{[\mathbf{y}]}$ is a positive-semidefinite matrix.*

In the original work of Lovász and Schrijver [23] it is shown that the cone $K_n(P)$ (even in the LS system) projected on $x_0 = 1$ is exactly the integral hull of the original LP relaxation, while one

can optimize over $K_t(P)$ in time $n^{O(t)}$, given that the original relaxation admits a (weak) polytime separation oracle. The algorithm in this section, as well as the one in Section 5, both rely heavily on the following facts, which follow easily from the above definition:

Fact 2.2. *For any vector $\mathbf{x} \in [0, 1]^n$ such that $(1, \mathbf{x}) \in K_t(P)$, and corresponding moment vector $\mathbf{y} \in \mathcal{P}_2$, and for any $i \in [n]$ such that $x_i > 0$, the rescaled column vector $\frac{1}{x_i} Y[\mathbf{y}] \mathbf{e}_i$ is in $K_{t-1}(P) \cap \{(1, \mathbf{x}') \mid \mathbf{x}' \in [0, 1]^n\}$.*

Fact 2.3. *For any vector $\mathbf{x} \in [0, 1]^n$ such that $(1, \mathbf{x}) \in K_t(P)$, and any coordinate j such that x_j is integral, for all $t' < t$, any vector \mathbf{x}' such that $(1, \mathbf{x}') \in K_{t'}(P)$ derived from \mathbf{x} by one or more steps as in Fact 2.2, we have $x'_j = x_j$.*

3 Two Approaches for Proving Theorem 1.1

In the following sections, we let (X, \mathcal{S}) denote a SET COVER instance with items X and cover-sets \mathcal{S} where each $S \in \mathcal{S}$ has cost $c(S)$. We use n to denote the number of items in X and m to denote the number of cover-sets in \mathcal{S} .

3.1 Sketch of a Combinatorial Proof

Recall that the standard greedy algorithm for approximating SET COVER iteratively selects the cover-set S of minimum density $c(S)/|S \setminus \bigcup_{T \in \mathcal{C}} T|$ where \mathcal{C} is the collection of cover-sets already chosen. The approximation guarantee of this algorithm is H_b , where b is the size of the largest cover-set. Our algorithm builds on this result simply by guessing up to d cover-sets in the optimal solution before running the greedy algorithm. However, some of the cover-sets in \mathcal{S} that were not guessed (and might still contain uncovered items) are discarded before running the greedy algorithm. Specifically, we discard the cover-sets that contain more than $\frac{n}{d}$ uncovered items after initially guessing the d sets. We show that for some choice of d sets, no remaining set in the optimum solution covers more than $\frac{n}{d}$ uncovered items. Thus, running the greedy algorithm on the remaining sets is actually an $H_{n/d}$ -approximation. The full description of the algorithm along with all details of the proof are in Appendix A

3.2 Proof Based on the Lovász-Schrijver System

In this section we provide an alternative LP-based approximation algorithm for SET COVER with the same performance as in Section 3.1, illustrating the power of the so-called lift-and-project systems. Consider the standard LP relaxation

$$\begin{aligned} \text{minimize} \quad & \sum_{S \in \mathcal{S}} c(S) x_S \\ \text{subject to} \quad & \sum_{S \ni i} x_S \geq 1 & \forall i \in X & (1) \\ & 0 \leq x_S \leq 1 & \forall S \in \mathcal{S} & (2) \end{aligned}$$

for SET COVER. Now, consider the corresponding feasibility LP where instead of explicitly minimizing the objective function, we add the following bound on the objective function as a constraint

(we will later guess the optimal value q by binary search):

$$\sum_{S \in \mathcal{S}} c(S)x_S \leq q. \quad (3)$$

We will work with the feasibility LP consisting only of constraints (1), (2) and (3) (and no objective function). Denote the corresponding polytope of feasible solutions by P_q .

In what follows we strengthen polytope P_q using the Lovász-Schrijver lift-and-project system. Next we show that the level- d Lovász-Schrijver relaxation $K_d(P_q)$ can give a $H_{\frac{n}{d}}$ -factor approximation algorithm. We note here that applying the Lovász-Schrijver system to the feasibility P_q (which includes the objective function as a constraint) and not on the standard LP relaxation of SET COVER is crucial, since by Alekhovich et al. [1] the latter LP has a very bad integrality gap even when strengthened by $\Omega(n)$ rounds of LS_+ (which is even stronger than LS).

To that end, let q be the smallest value such that the level- d LS tightening of P_q is not empty (note that $q \leq \text{OPT}$). The value q can be found through binary search (note that in each stage of the binary search we attempt to check $K_d(P_{q'})$ for emptiness for some q' , which takes time $m^{O(d)}$). Our goal is to show that for this q we can find a SET COVER of cost at most $q \cdot H_{\frac{n}{d}}$.

Let $\mathbf{x}^{(d)}$ be such that $(1, \mathbf{x}^{(d)}) \in K_d(P_q)$. For any coordinate i in the support of $\mathbf{x}^{(d)}$ we can invoke Fact 2.2 and get a vector $\mathbf{x}^{(d-1)}$ such that $(1, \mathbf{x}^{(d-1)}) \in K_{d-1}(P_q)$ and $\mathbf{x}_i^{(d-1)} = 1$. By Fact 2.3, by iterating this step, we eventually obtain a vector $\mathbf{x}^{(0)} \in P_q$ which is integral in at least d coordinates. Note that by constraint (3), this solution has cost at most q . We refer to this subroutine as the *Conditioning Phase*, which is realized in d many inductive steps.

If at some level $0 \leq i \leq d$, the sets whose coordinates in $\mathbf{x}^{(i)}$ are set to 1 cover all universe elements X , we have solved the SET COVER instance with cost

$$\sum_{S: \mathbf{x}_S^{(i)} = 1} C(S) \leq q \leq \text{OPT}.$$

Otherwise, we need to solve a smaller instance of SET COVER defined by all elements $Y \subseteq X$ not already covered, using cover-sets $\mathcal{T} = \{S \cap Y \mid \mathbf{x}_S^{(0)} > 0\}$. We introduce some structure in the resulting instance (Y, \mathcal{T}) of SET COVER by choosing the indices we condition on as in the proof of Lemma A.2. This gives the following Lemma, whose proof is similar to that of Lemma A.2.

Lemma 3.1. *If at each step of the Conditioning Phase we choose the set S in the support of the current solution $\mathbf{x}^{(d')}$ containing the most uncovered elements in X , then for all $T \in \mathcal{T}$ we have $|T| \leq \frac{n}{d}$.*

Proof. For $1 \leq i \leq d$ let S_i denote the cover-set chosen in the Conditioning Phase for level i and for $0 \leq i \leq d$ let $\mathcal{C}_i = \{S_d, S_{d-1}, \dots, S_{i+1}\}$ (with $\mathcal{C}_d = \emptyset$). We show that at every iteration $1 \leq i \leq d$ we have $|S \setminus \bigcup_{T \in \mathcal{C}_i} T| \leq \frac{n}{d-i+1}$ for every $S \in \mathcal{S} \setminus \mathcal{C}_i$ with $\mathbf{x}_S^{(i)} > 0$.

For $1 \leq i \leq d$, let $\alpha_i := |S_i \setminus \bigcup_{T \in \mathcal{C}_i} T|$. Since we chose the largest (with respect to the uncovered items) cover-set S_i in the support of $\mathbf{x}^{(i)}$ we have $|S' \setminus \bigcup_{T \in \mathcal{C}_i} T| \leq \alpha_i$ for every cover-set $S' \in \mathcal{S} \setminus \mathcal{C}_i$ with $\mathbf{x}_{S'}^{(i)} > 0$. For $2 \leq i \leq d$ we also have $\alpha_{i-1} \leq \alpha_i$ because $\mathcal{C}_i \subseteq \mathcal{C}_{i-1}$ and because we chose S_i instead of S_{i-1} in the Conditioning Phase for level i .

So, $\alpha_d \geq \alpha_{d-1} \geq \dots \geq \alpha_1$. Now, each item j covered by \mathcal{C}_0 contributes 1 to α_i for the earliest index i for which $j \in S_i$, so $\sum_{i=1}^d \alpha_i \leq n$. This implies that $\alpha_i \leq \frac{n}{d-i+1}$ for each $1 \leq i \leq d$. Therefore, every set in the support of $\mathbf{x}^{(0)}$ has at most $\frac{n}{d}$ elements that are not already covered by \mathcal{C}_0 . So, the instance (Y, \mathcal{T}) has $|T| \leq \frac{n}{d}$ for any $T \in \mathcal{T}$. \square

Let \mathcal{D} be the collection of cover-sets chosen as in Lemma 3.1. Observe that the vector $\mathbf{x}^{(0)}$ projected on the cover-sets $\mathcal{S} \setminus \mathcal{D}$ that were *not* chosen in the Conditioning Phase is feasible for the LP relaxation of the instance (Y, \mathcal{T}) . In particular, the cost of the LP is at most $q - \sum_{S \in \mathcal{D}} c(S)$, and by Lemma 3.1 all cover-sets have size at most $\frac{n}{d}$. By Theorem A.1, the greedy algorithm will find a solution for (Y, \mathcal{T}) of cost at most $H_{\frac{n}{d}} \cdot (q - \sum_{S \in \mathcal{D}} c(S))$. Altogether, this gives a feasible solution for (X, \mathcal{S}) of cost

$$H_{\frac{n}{d}} \cdot \left(q - \sum_{S \in \mathcal{D}} c(S) \right) + \sum_{S \in \mathcal{D}} c(S) \leq H_{\frac{n}{d}} \cdot q \leq H_{\frac{n}{d}} \cdot \text{OPT}.$$

4 Linear Sherali-Adams Integrality Gap for Set Cover

The level- ℓ Sherali-Adams relaxation is a tightened LP that can be derived systematically starting with any 0-1 LP relaxation. While in this work we are interested in tightening the SET COVER polytope, the process we describe below is applicable to any other relaxation.

Definition 4.1 (The Sherali-Adams system). *Consider a polytope over the variables y_1, \dots, y_n defined by finitely many constraints (including the box-constraints $0 \leq y_i \leq 1$). The level- ℓ Sherali-Adams relaxation is an LP over the variables $\{y_A\}$ where A is any subset of $\{1, 2, \dots, n\}$ of size at most $\ell + 1$, and where $y_\emptyset = 1$. For every constraint $\sum_{i=1}^n a_i y_i \geq b$ of the original polytope and for every disjoint $P, E \subseteq \{1, \dots, n\}$ with $|P| + |E| \leq \ell$, the following is a constraint of the level- ℓ Sherali-Adams relaxation*

$$\sum_{i=1}^n a_i \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T \cup \{i\}} \geq b \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T}.$$

We will prove Theorem 1.2 in this section. For this we will need two ingredients: (a) appropriate instances, and (b) a solution of the Sherali-Adams LP as described in Definition 4.1. Our hard instances are described in the following lemma, which is due to Alekhovich et al. [1].

Lemma 4.2 (SET COVER instances with no small feasible solutions).

For every $\varepsilon > \eta > 0$, and for all sufficiently large n , there exist SET COVER instances over a universe of n elements and n cover-sets, such that:

- (i) *Every element of the universe appears in exactly $(\varepsilon - \eta)n$ cover-sets, and*
- (ii) *There is no feasible solution that uses less than $\log_{1+\varepsilon} n$ cover-sets.*

In order to prove Theorem 1.2 we will invoke Lemma 4.2 with appropriate parameters. Then we will define a vector solution for the level- ℓ Sherali-Adams relaxation as described.

Lemma 4.3. *Consider a SET COVER instance on n cover-sets as described in Lemma 4.2. Let f denote the number of cover-sets covering every element of the universe. For $f \geq 3\ell$, the vector \mathbf{y} indexed by subsets of $\{1, \dots, n\}$ of size at most $\ell + 1$ defined as $y_A := \frac{(f-\ell-1)!}{(f-\ell-1+|A|)!}$, $\forall A \subseteq \{1, \dots, n\}$, $|A| \leq \ell + 1$, satisfies the level- ℓ Sherali-Adams LP relaxation of the SET COVER polytope.*

The proof of Lemma 4.3 involves a number of extensive calculations which we give in Section 4.1. Assuming the lemma, we are ready to prove Theorem 1.2.

Proof of Theorem 1.2. Fix $\varepsilon > 0$ and invoke Lemma 4.2 with $\eta = \varepsilon^2$ to obtain a SET COVER instance on n universe elements and n cover-sets for which (i) every universe element is covered by exactly $(\varepsilon - \varepsilon^2)n$ cover-sets, and (ii) no feasible solution exists of cost less than $\log_{1+\varepsilon} n$. Note that in particular (i) implies that in the SET COVER LP relaxation, every constraint has support exactly $f = (\varepsilon - \varepsilon^2)n$.

Set $\ell = \frac{\gamma(\varepsilon - \varepsilon^2)}{1+\gamma}n$ and note that $f/\ell \geq 3$, since $\gamma \leq \frac{1}{2}$. This means we can define a feasible level- ℓ Sherali-Adams solution as described in Lemma 4.3. The values of the singleton variables are set to

$$y_{\{i\}} = \frac{1}{(\varepsilon - \varepsilon^2)n - \ell} = \frac{1 + \gamma}{(\varepsilon - \varepsilon^2)n}.$$

But then, the integrality gap is at least

$$\frac{\text{OPT}}{\sum_{i=1}^n y_{\{i\}}} \geq \frac{\varepsilon - \varepsilon^2}{1 + \gamma} \cdot \log_{1+\varepsilon} n = \frac{\varepsilon - \varepsilon^2}{(1 + \gamma) \ln(1 + \varepsilon)} \ln n.$$

The lemma follows once we observe that $\ln(1 + \varepsilon) = \varepsilon - \frac{1}{2}\varepsilon^2 + \Theta(\varepsilon^3)$. \square

4.1 Proof of Lemma 4.3

Recall that f denotes the support of every cover constraint in the SET COVER relaxation (i.e. the number of cover-sets every element belongs to in the instance described in Lemma 4.2). Also recall that

$$y_A := \frac{(f - \ell - 1)!}{(f - \ell - 1 + |A|)!}, \quad \forall A \subseteq \{1, \dots, n\}, \quad |A| \leq \ell + 1.$$

To show feasibility of the above vectors, we need to study two types of constraints, i.e. the so called box-constraints

$$0 \leq \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T} \leq 1, \quad \forall P, E \subseteq \{1, \dots, n\}, \quad |P| + |E| \leq \ell + 1, \quad (4)$$

as well as the covering constraints

$$\sum_{i \in D} \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T \cup \{i\}} \geq \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T}, \quad \forall P, E \subseteq \{1, \dots, n\}, \quad |P| + |E| \leq \ell, \quad (5)$$

where $D \subseteq \{1, \dots, n\}$ is a set of f many cover-sets covering some element of the universe. The symmetry of the proposed solutions allows us to significantly simplify the above expressions, by noting that for $|P| = p$ and $|E| = e$ we have

$$\sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T} = \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{(f - \ell - 1)!}{(f - \ell - 1 + p + t)!}$$

For the sake of exposition, we show that our LP solution satisfies the two different kinds of constraints in two different lemmata. Set $x = f - \ell - 1 + p$, and note that if $f \geq 3\ell$, then since $e \leq \ell$ we have $x \geq \frac{5}{3}e$. Thus, the feasibility of the box constraints (4) (for our solution) is implied by the following combinatorial lemma.

Lemma 4.4. For $x \geq \frac{5}{3}e$ we have

$$0 \leq \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(x+t)!} \leq \frac{1}{(x-p)!}$$

Proof. First we show the lower bound, for which we study two consecutive summands. We note that

$$\begin{aligned} \binom{e}{2t} \frac{1}{(x+2t)!} - \binom{e}{2t+1} \frac{1}{(x+2t+1)!} &= \binom{e}{2t} \frac{1}{(x+2t)!} \left(1 - \frac{\binom{e}{2t+1}}{\binom{e}{2t}} \frac{1}{x+2t+1} \right) \\ &= \binom{e}{2t} \frac{1}{(x+2t)!} \left(1 - \frac{e-2t}{2t+1} \frac{1}{x+2t+1} \right) \\ &\geq \binom{e}{2t} \frac{1}{(x+2t)!} \left(1 - \frac{e}{2t+1} \frac{1}{x+2t+1} \right) \\ &\geq \binom{e}{2t} \frac{1}{(x+2t)!} \left(1 - \frac{e}{x} \right). \end{aligned}$$

Since $x \geq e$, every two consecutive summands add up to a non-negative value. Thus, if the number of summands is even (that is, if e is odd), the lower bound follows, and if the number of summands is odd, the bound also follows, since for even e , the last summand is positive.

Now we show the upper bound. Note that, since p does not appear in the sum, it suffices to bound the sum by the smaller value $\frac{1}{x!}$. This is facilitated by noting that

$$\sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(x+t)!} = \frac{1}{x!} \sum_{t=0}^e (-1)^t \frac{1}{t!} \frac{\binom{x+e}{x+t}}{\binom{x+e}{e}}.$$

Hence, it suffices to show that

$$\sum_{t=0}^e (-1)^t \frac{1}{t!} \binom{x+e}{x+t} \leq \binom{x+e}{e}. \quad (6)$$

As before, we analyze the sum of two consecutive terms (this time in the above sum). This is done in the next claim.

Claim 4.5. For $x \geq \frac{5}{3}e$ we have

$$\frac{1}{(2t)!} \binom{x+e}{x+2t} - \frac{1}{(2t+1)!} \binom{x+e}{x+2t+1} \leq \binom{x+e}{x+2t} - \binom{x+e}{x+2t+1}$$

Proof. We divide both sides of the desired inequality by $\binom{x+e}{x+2t}$ to obtain the equivalent statement

$$\begin{aligned} \frac{1}{(2t)!} - \frac{1}{(2t+1)!} \frac{e-2t}{x+2t+1} &\leq 1 - \frac{e-2t}{x+2t+1} \\ \Leftrightarrow \frac{e-2t}{x+2t+1} \left(1 - \frac{1}{(2t+1)!} \right) &\leq 1 - \frac{1}{(2t)!} \end{aligned}$$

Note that the above is tight for $t = 0$. For $t > 0$, and since $x \geq \frac{5}{3}e$, we have $\frac{e-2t}{x+2t+1} < \frac{3}{5}$ which is small enough to compensate for the worst-case ratio of the expressions involving factorials, which occurs for $t = 1$. \square

Continuing our proof of Lemma 4.4, first suppose that e is odd. Then Claim 4.5 implies that

$$\sum_{t=0}^e (-1)^t \frac{1}{t!} \binom{x+e}{x+t} \leq \sum_{t=0}^e (-1)^t \binom{x+e}{x+t} = \binom{x+e-1}{e} \leq \binom{x+e}{e},$$

which gives the required condition (6) for odd e . If e is a positive even integer, then again Claim 4.5 implies that

$$\sum_{t=0}^e (-1)^t \frac{1}{t!} \binom{x+e}{x+t} \leq \sum_{t=0}^{e-1} (-1)^t \binom{x+e}{x+t} + \frac{1}{e!} = \binom{x+e-2}{e} + \frac{1}{e!} < \binom{x+e}{e}.$$

Finally, if $e = 0$ then (6) holds with equality. This concludes the proof. \square

Now we turn our attention to cover constraints (5). Recall that the values y_A only depend on the size of A . Since f and ℓ are fixed in the context of Lemma 4.3, for $|P| = p$ and $|E| = e$ we can define

$$H_{e,p} := \sum_{\emptyset \subseteq T \subseteq E} (-1)^{|T|} y_{P \cup T} = \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{(f-\ell-1)!}{(f-\ell-1+p+t)!}.$$

Thus the left-hand-side of (5) (for fixed P and E) involves only expressions of the form $H_{e,p}$, $H_{e,p+1}$, or 0, depending on the relationship between the sets P, E and $\{i\}$, while the right-hand-side equals $H_{e,p}$.

More concretely, let $|D \cap P| = p_1$, $|D \cap E| = e_1$, $|P \setminus D| = p_0$ and $|E \setminus D| = e_0$, where $p_0 + p_1 = |P| = p$ and $e_0 + e_1 = |E| = e$, and recall that E and P are disjoint. Then observe that the p_1 indices in $D \cap P$ each contribute $H_{e,p}$ to the left-hand-side of (5). In addition, the e_1 indices $i \in D \cap E$ each contribute 0 to the left-hand-side. This follows because each $T \subseteq E \setminus \{i\}$ can be paired with $T' := T \cup \{i\} \subseteq E$ and the terms $y_{P \cup T \cup \{i\}}$ and $y_{P \cup T' \cup \{i\}}$ are identical, while they appear in the sum with opposite signs. Finally, the remaining $f - p_1 - e_1$ indices contribute each $H_{e,p+1}$. Overall, for our proposed solution, Constraint (5) can be rewritten as

$$(f - p_1 - e_1)H_{e,p+1} + p_1 H_{e,p} \geq H_{e,p}$$

Clearly, for $p_1 > 0$, the above constraint is satisfied. Hence, we may assume that $|P \cap D| = \emptyset$, and so $|P| = p_1 = p$. Note also that the value e_1 does not affect $H_{e,p}$, thus the above inequality holds for all $e_1 \leq e$ iff it holds for $e_1 = e$ and $e_0 = 0$. To summarize, to show that Constraint (5) is satisfied, we need only show that

$$(f - e)H_{e,p+1} \geq H_{e,p} \tag{7}$$

for $H_{e,p} = \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{(f-\ell-1)!}{(f-\ell-1+p+t)!}$, and $e + p \leq \ell$. Note that the value $(f - \ell - 1)!$ in the numerator appears in both sides. Also $f \geq 3\ell$ and $e \leq \ell$ implies that $e < 2\ell \leq f - \ell$. Finally recall that in (5) we have $p + e = |P| + |L| \leq \ell$. Hence, to show that (7) (and thus Constraint (5)) is satisfied, it remains only to show the following lemma.

Lemma 4.6. *For $e < f - \ell$ and $p + e \leq \ell$ we have*

$$(f - e) \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(f - \ell + p + t)!} - \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(f - \ell - 1 + p + t)!} \geq 0$$

Proof. We rewrite the left-hand-side as

$$\begin{aligned}
& \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(f-\ell-1+p+t)!} \left(\frac{f-e}{f-\ell+p+t} - 1 \right) \\
&= \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(f-\ell-1+p+t)!} \frac{\ell-(p+e)-t}{f-\ell+p+t} \\
&= (\ell-(p+e)) \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{1}{(f-\ell+p+t)!} - \sum_{t=0}^e (-1)^t \binom{e}{t} \frac{t}{(f-\ell+p+t)!}
\end{aligned}$$

The first sum of the last expression is non-negative, as it is a multiple of the box-constraints we have already proven. So is its coefficient, since $p+e \leq \ell$. Note that for $e=0$ and $p=\ell$ the above expression equals 0 (and in particular Constraint (5) is tight). Next we show that for other values of p, e (either $p < \ell$ or $e > 0$) the above expression is strictly positive. For this it suffices to show that $\sum_{t=0}^e (-1)^t \binom{e}{t} \frac{t}{(f-\ell+p+t)!} \leq 0$.

We proceed again by analyzing every two consecutive terms. We observe that for $t \geq 1$ we have

$$\begin{aligned}
& -\binom{e}{2t-1} \frac{2t-1}{(f-\ell+p+2t-1)!} + \binom{e}{2t} \frac{2t}{(f-\ell+p+2t)!} \\
&= \frac{\binom{e}{2t-1}}{(f-\ell+p+2t-1)!} \left(-(2t-1) + \frac{e-2t+1}{f-\ell+p+2t} \right) \\
&< \frac{\binom{e}{2t-1}}{(f-\ell+p+2t-1)!} \left(-(2t-1) + \frac{e}{f-\ell} \right) \\
&< \frac{\binom{e}{2t-1}}{(f-\ell+p+2t-1)!} \left(-1 + \frac{e}{f-\ell} \right)
\end{aligned}$$

which is non-positive, since $f-\ell \geq e$. This argument shows that every two consecutive summands of $\sum_{t=0}^e (-1)^t \binom{e}{t} \frac{t}{(f-\ell+p+t)!}$ add up to a non-positive value. If e is even, then we are done, while for odd e the unmatched summand (for $t=e$) is negative, and so the lemma follows. \square

5 An LS_+ -based PTAS for Knapsack

We consider the KNAPSACK problem: We are given n items which we identify with the integers $[n]$, and each item $i \in [n]$ has some associated (nonnegative) reward r_i and cost (or size) c_i . The goal is to choose a set of items which fit in the knapsack, i.e. whose total cost does not exceed some bound C , so as to maximize the total reward. In what follows we will use the LP $\{\max \sum_{i=1}^n r_i x_i : \sum_{i=1}^n c_i x_i \leq C \text{ \& } 0 \leq x_i \leq 1 \forall i \in [n]\}$, which is the natural relaxation for KNAPSACK.

Denote the polytope associated with the KNAPSACK LP relaxation by P . We will consider the SDP derived by applying sufficiently many levels of LS_+ (as defined in Section 3.2) to the above LP. That is, for some $\ell > 0$, we consider the SDP

$$\begin{aligned}
& \text{maximize} && \sum_{i=1}^n r_i x_i \\
& \text{subject to} && (1, \mathbf{x}) \in K_\ell^+(P).
\end{aligned}$$

There is a well-known simple greedy algorithm for KNAPSACK: Sort the items by decreasing order of r_i/c_i , and add them to the knapsack one at a time until the current item does not fit. The following lemma (which is folklore) relates the performance of the greedy algorithm to the value of the LP relaxation P :

Lemma 5.1. *Let \mathbf{x} be a solution to P , and R_G be the reward given by the greedy algorithm. Then $\sum_i r_i x_i \leq R_G + \max_i r_i$.*

This gives a trivial bound of 2 on the integrality gap, assuming that $c_i \leq C$ for all i (that is, that each item can be a solution on its own), since we then have $R_G + \max_i r_i \leq 2\text{OPT}$.¹ We will see later that the above assumption can essentially be enforced, that is, that we can ignore items with reward greater than C (see Lemma 5.6). Lemma 5.1 has the following easy corollary: Consider the above greedy algorithm, with the modification that we first add all items which have $x_i = 1$ and discard all items which have $x_i = 0$. Then the following holds:

Corollary 5.2. *Let \mathbf{x} be a solution to P , and let R'_G be the total reward given by the above modified greedy algorithm. Then $\sum_i r_i x_i \leq R'_G + \max_{i: 0 < x_i < 1} r_i$.*

We will show that, for any constant $\varepsilon > 0$, there is a constant L_ε such that the SDP relaxation for KNAPSACK arising from level L_ε of LS_+ has integrality gap at most $1 + O(\varepsilon)$. For the Lasserre hierarchy, this has been shown for level $1/\varepsilon$ [19]. We will show this for $L_\varepsilon = 1/\varepsilon^3$ in the case of LS_+ .

Our rounding algorithm will take as input the values of the KNAPSACK instance $(r_i)_i, (c_i)_i$, and C , an optimal solution \mathbf{x} s.t. $(1, \mathbf{x}) \in K_\ell^+(P)$ (for some level $\ell > 0$, initially $\ell = L_\varepsilon$), and parameters ε and ρ . The parameter ρ is intended to be the threshold $\varepsilon \cdot \text{OPT}$ in the set $S_{\varepsilon \cdot \text{OPT}} = \{i \mid r_i > \varepsilon \cdot \text{OPT}\}$. Rather than guessing a value for OPT , though, we will simply try all values of $\rho \in \{r_i \mid i \in [n]\} \cup \{0\}$ and note that for exactly one of those values, the set $S_{\varepsilon \cdot \text{OPT}}$ coincides with the set $\{i \mid r_i > \rho\}$ (also note that ρ is a parameter of the rounding, and not involved at all in the SDP relaxation).

The intuition behind our rounding algorithm is as follows: As we did for SET COVER, we would like to repeatedly “condition” on setting some variable to 1, by using Fact 2.2. If we condition only on (variables corresponding to) items in $S_{\varepsilon \cdot \text{OPT}}$, then after at most $1/\varepsilon$ iterations, the SDP solution will be integral on that set, and then by Corollary 5.2 the modified greedy algorithm will give a $1 + O(\varepsilon \cdot \text{OPT})$ approximation relative to the value of the objective function (since items outside $S_{\varepsilon \cdot \text{OPT}}$ have reward at most $\varepsilon \cdot \text{OPT}$). The problem with this approach (and the reason why LP hierarchies do not work), is the same problem as for SET COVER: the conditioning step does not preserve the value of the objective function. While the optimum value of the SDP is at least OPT by definition, after conditioning, the value of the new solution may be much smaller than OPT , which then makes the use of Corollary 5.2 meaningless. The key observation is that the use of SDPs ensures that we can choose some item to condition on without any decrease in the objective function (see Lemma 5.4).² A more refined analysis shows that we will be able condition *specifically on items in $S_{\varepsilon \cdot \text{OPT}}$* without losing too much. Counter-intuitively, we then need to show that the algorithm does not perform an unbounded number of conditioning steps which *increase* the objective value (see Lemma 5.8). Our rounding algorithm **KS-Round** is described in Algorithm 1, while the performance guarantee is described in Section 5.1.

¹Here, as before, OPT denotes the optimal 0-1 solution.

²Note that this is crucial for a maximization problem like KNAPSACK, while for a minimization problem like SET COVER it does not seem helpful (and indeed, by the integrality gap of Alekhovich et al. [1], we know it does not help).

Algorithm 1 $\text{KS-Round}((r_i)_i, (c_i)_i, C, \mathbf{x}, \varepsilon, \rho)$

- 1: Let $\mathbf{y} \in \mathcal{P}_2$ be the moment vector associated with $(1, \mathbf{x})$.
 - 2: Let $S_\rho \leftarrow \{i \mid r_i > \rho\}$, and let $S^b \leftarrow \{i \mid x_i = b\}$ for $b = 0, 1$.
 - 3: **if** $S_\rho \subseteq S^0 \cup S^1$ **then**
 - 4: Run the modified greedy algorithm.
 - 5: **else if** $\sum_{i \in S_\rho \setminus S^1} r_i x_i < \varepsilon \cdot \sum_{i=1}^n r_i x_i$ **then**
 - 6: Run the modified greedy algorithm on items in $([n] \setminus S_\rho) \cup S^1$.
 - 7: **else if** there is some $i \in S_\rho \setminus (S^0 \cup S^1)$ s.t. $\sum_{j=1}^n r_j y_{\{i,j\}} \geq (1 - \varepsilon^2) x_i \cdot \sum_{j=1}^n r_j x_j$ **then**
 - 8: Run $\text{KS-Round}((r_i)_i, (c_i)_i, C, \frac{1}{x_i} Y^{[\mathbf{y}]} \mathbf{e}_i, \varepsilon, \rho)$. ▷ See Fact 2.2
 - 9: **else**
 - 10: Choose $i \in [n] \setminus (S_\rho \cup S^0)$ s.t. $\sum_{j=1}^n r_j y_{\{i,j\}} > (1 + \varepsilon^3) x_i \cdot \sum_{j=1}^n r_j x_j$ ▷ See Lemma 5.5
 - 11: Run $\text{KS-Round}((r_i)_i, (c_i)_i, C, \frac{1}{x_i} Y^{[\mathbf{y}]} \mathbf{e}_i, \varepsilon, \rho)$. ▷ See Fact 2.2
 - 12: **end if**
-

5.1 Analysis of Algorithm 1

Before we analyze the performance guarantee of the algorithm, there is one more simple fact about both LS and LS_+ which we will use in this section.

Fact 5.3. *Given a solution $(1, \mathbf{x}) \in K_\ell(P)$ for some $\ell \geq 1$ and corresponding moment vector $\mathbf{y} \in \mathcal{P}_2$, if $x_i = 1$ for some $i \in [n]$, then $Y^{[\mathbf{y}]} \mathbf{e}_i = Y^{[\mathbf{y}]} \mathbf{e}_0 (= (1, \mathbf{x}))$.*

This fact follows easily from the fact that $Y^{[\mathbf{y}]} (\mathbf{e}_0 - \mathbf{e}_i) \in K_{\ell-1}(P)$, and since if $x_i = y_{\{i\}} = 1$ then the first entry in the above vector is 0, which for the conified polytope $K_{\ell-1}(P)$ only holds for the all-zero vector.

Now, consider the set S_ρ defined in Step 2. As we have pointed out, for the appropriate choice of ρ , this set coincides with the set $S_{\varepsilon\text{OPT}}$. Note that, by Corollary 5.2, if all the x_i values in S_ρ are integral, then Step 4 returns a solution with value R'_G satisfying

$$\text{OPT} \geq R'_G \geq \sum_i r_i x_i - \max_{i \notin S_\rho} r_i \geq \sum_i r_i x_i - \varepsilon \cdot \text{OPT}.$$

Now, if \mathbf{x} is the original SDP solution given to the rounding algorithm, this gives an upper bound of $1 + \varepsilon$ on the integrality gap (as well as a $(1 + \varepsilon)$ -approximation). However, in Steps 8 and 11, we recurse with a new SDP solution (to a lower level in the LS_+ hierarchy). Thus, our goal will be to arrive at an SDP solution which is integral on S_ρ (or assigns so little weight to S_ρ that we can ignore it), but we need to show that the objective function does not decrease too much during these recursive steps. To show this, we will rely crucially on the following easy lemma, which is also the only place where we use positive-semidefiniteness.

Lemma 5.4. *Let $(1, \mathbf{x})$ be a solution to $K_\ell^+(P)$ for some $\ell \geq 1$, with the corresponding moment*

vector \mathbf{y} . Then the solution satisfies

$$\sum_{i=1}^n r_i \sum_{j=1}^n r_j y_{\{i,j\}} \geq \left(\sum_{i=1}^n r_i x_i \right)^2 \left(= \sum_{i=1}^n r_i \sum_{j=1}^n x_i r_j x_j \right).$$

Proof. By the positive semidefiniteness of the moment matrix $Y^{[\mathbf{y}]}$, we have $\mathbf{a}^\top Y^{[\mathbf{y}]} \mathbf{a} \geq 0$ for any vector $\mathbf{a} \in \mathbb{R}^{n+1}$. Then the lemma follows immediately from this inequality, by letting $\mathbf{a} = (a_i)_i$ be the vector defined by $a_i = r_i$ for $i \in [n]$ and $a_0 = -\sum_{i=1}^n r_i x_i$. \square

Thus, there is some item $i \in [n]$ on which we can condition (by taking the new solution $\frac{1}{x_i} Y^{[\mathbf{y}]} \mathbf{e}_i$, without any decrease in the value of the objective function. Moreover, using the above lemma, we can now show that the algorithm is well-defined (assuming we have start with an SDP at a sufficiently high level of LS_+ for all the recursive steps).

Lemma 5.5. *If Step 10 is reached, then there exists an item $i \in [n] \setminus (S_\rho \cup S^0)$ satisfying*

$$\sum_{j=1}^n r_j y_{\{i,j\}} > (1 + \varepsilon^3) x_i \cdot \sum_{j=1}^n r_j x_j. \quad (8)$$

Proof. Note that if Step 10 is reached then we must have

$$\sum_{i \in S_\rho \setminus S^1} r_i x_i \geq \varepsilon \cdot \sum_{i=1}^n r_i x_i. \quad (9)$$

Moreover, for all $i \in S_\rho \setminus (S^0 \cup S^1)$ we have $\sum_{j=1}^n r_j y_{\{i,j\}} < (1 - \varepsilon^2) x_i \cdot \sum_{j=1}^n r_j x_j$. In particular, we have

$$\sum_{i \in S_\rho \setminus S^1} r_i \sum_{j=1}^n r_j y_{\{i,j\}} < (1 - \varepsilon^2) \sum_{i \in S_\rho \setminus S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j. \quad (10)$$

Thus (noting that all $i \in S^0$ contribute nothing to the following sums), we have

$$\begin{aligned}
\sum_{i \in ([n] \setminus (S_\rho \cup S^0)) \cup S^1} r_i \sum_{j=1}^n r_j y_{\{i,j\}} &= \sum_{i=1}^n r_i \sum_{j=1}^n r_j y_{\{i,j\}} - \sum_{i \in S_\rho \setminus S^1} r_i \sum_{j=1}^n r_j y_{\{i,j\}} \\
&\geq \left(\sum_{i=1}^n r_i x_i \right)^2 - \sum_{i \in S_\rho \setminus S^1} r_i \sum_{j=1}^n r_j y_{\{i,j\}} && \text{by Lemma 5.4} \\
&> \left(\sum_{i=1}^n r_i x_i \right)^2 - (1 - \varepsilon^2) \sum_{i \in S_\rho \setminus S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j && \text{by (10)} \\
&= \sum_{i \in ([n] \setminus S_\rho) \cup S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j + \varepsilon^2 \sum_{i \in S_\rho \setminus S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j \\
&\geq \sum_{i \in ([n] \setminus S_\rho) \cup S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j + \varepsilon^3 \sum_{i=1}^n r_i x_i \cdot \sum_{j=1}^n r_j x_j && \text{by (9)} \\
&\geq (1 + \varepsilon^3) \cdot \sum_{i \in ([n] \setminus S_\rho) \cup S^1} r_i x_i \cdot \sum_{j=1}^n r_j x_j.
\end{aligned}$$

Therefore, there is some $i \in ([n] \setminus (S_\rho \cup S^0)) \cup S^1$ satisfying (8). We only need to show that $i \notin S^1$, that is, that $x_i < 1$. However, by Fact 5.3, if $x_i = 1$ then $y_{\{i,j\}} = x_j$, making inequality (8) impossible for such i . \square

Now that we have shown the algorithm to be well-defined, let us start to bound the depth of the recursion. We will use the following lemma to show that after a bounded number of recursive calls in Step 8, the SDP solution becomes integral on S_ρ :

Lemma 5.6. *Let $(1, \mathbf{x})$ be a solution to K_ℓ^+ for some $\ell \geq 1$. Then for all items $i \in [n] \setminus S^1$ such that $c_i > C - \sum_{j \in S^1} c_j$, we have $x_i = 0$.*

Proof. Suppose, for the sake of contradiction, that there is some item i satisfying this property, with $0 < x_i < 1$. By Fact 5.3, for every $j \in S^1$ we have $y_{\{i,j\}} = x_i$. However, by Fact 2.2, the vector $\frac{1}{x_i} Y^{[\mathbf{y}]} \mathbf{e}_i$ satisfies the constraints of $K_{\ell-1}^+$, and in particular the capacity constraint. But this is a contradiction, since we have

$$\frac{1}{x_i} \sum_{j=1}^n c_j y_{\{i,j\}} = c_i + \sum_{j \in S^1} c_j + \sum_{j \in [n] \setminus (S^1 \cup \{i\})} c_j x_j > C.$$

\square

Since no $1/\varepsilon$ items from $S_{\varepsilon\text{OPT}}$ can fit simultaneously in the knapsack, we have the following:

Corollary 5.7. *For ρ s.t. $S_\rho = S_{\varepsilon\text{OPT}}$, the recursion in Step 8 cannot be repeated $\lceil 1/\varepsilon \rceil$ times. Moreover, if this step is repeated $(\lceil 1/\varepsilon \rceil - 1)$ times, then after these recursive calls, we have $x_i \in \{0, 1\}$ for all $i \in S_\rho$.*

We can now bound the total depth of the recursion in the algorithm.

Lemma 5.8. *For all positive $\varepsilon \leq 0.3$, the algorithm performs the recursion in Step 11 less than $\lceil 1/\varepsilon^3 \rceil$ times.*

Proof. This follows by tracking the changes to the value of the objective function. Each time Step 8 is performed, the value of the objective function changes by a factor at least $(1 - \varepsilon^2)$, while each time Step 11 is performed, this value changes by a factor greater than $(1 + \varepsilon^3)$. Assuming, for the sake of contradiction, that Step 11 is performed $\lceil 1/\varepsilon^3 \rceil$ times, then by Corollary 5.7 the total change is at least a factor $(1 - \varepsilon^2)^{\lceil 1/\varepsilon \rceil - 1} (1 + \varepsilon^3)^{\lceil 1/\varepsilon^3 \rceil} > 2$. However, this is a contradiction, since initially the value of the objective function is at least OPT , and the integrality gap is always at most 2. \square

Finally, we can prove that the algorithm gives a $(1 + O(\varepsilon))$ -approximation, thus bounding the integrality gap. Theorem 1.3 follows immediately from the following.

Theorem 5.9. *For sufficiently small $\varepsilon > 0$, given an optimal (maximum feasible) solution $(1, \mathbf{x})$ to $K_\ell^+(P)$ for $\ell = O(1/\varepsilon^3)$, there is a value $\rho \in \{r_i \mid i \in [n]\} \cup \{0\}$ such that algorithm **KS-Round** finds a solution to the KNAPSACK instance with total value (reward) at least $(1 - O(\varepsilon)) \cdot \sum_{i=1}^n r_i x_i$.*

Proof of Theorem 5.9. It is easy to see that for (exactly one) $\rho \in \{r_i \mid i \in [n]\} \cup \{0\}$ we have $S_\rho = S_{\varepsilon \text{OPT}}$. Choose this value of ρ .

Note that the only time the value of the objective function can decrease is in Step 8. Let Φ be the initial value of the objective function. Then by Corollary 5.7, at the end of the recursion, the new value of the objective function will be at least $(1 - \varepsilon^2)^{1/\varepsilon} \cdot \Phi = (1 - O(\varepsilon))\Phi$. Thus, it suffices to show the theorem relative to the final value of the objective function (as opposed to Φ)³.

Let us consider the algorithm step-by-step. If the algorithm terminates in Step 4, then since for all $i \in [n] \setminus S_{\varepsilon \text{OPT}}$ we have $r_i \leq \varepsilon \cdot \text{OPT}$, by Corollary 5.2 the rounding loses at most a $(1 + O(\varepsilon))$ factor relative to the value of the objective function.

If the algorithm terminates in Step 6, then similarly, the rounding loses at most a $(1 + O(\varepsilon))$ -factor relative to $\sum_{i \in [n] \setminus S_\rho} r_i x_i > (1 - \varepsilon) \sum_{i=1}^n r_i x_i$.

Finally, note that by Corollary 5.7 and Lemma 5.8, we have sufficiently many levels of the hierarchy to justify the recursive Steps 8 and 11, and note that the choice of item i in Step 10 is well-defined by Lemma 5.4. \square

6 Conclusion

The known sub-exponential time approximation algorithms place SET COVER in a distinct category from other optimization problems like MAX-3SAT in the following sense. Though one can achieve provably hard approximation factors for SET COVER in sub-exponential time, Moshkovitz and Raz [25] show that improving over the easy $\frac{8}{7}$ -approximation for MAX-3SAT [18] by any constant requires time $2^{n^{1-o(1)}}$, assuming the ETH.

Rather, SET COVER lies in the same category of problems as CHROMATIC NUMBER and CLIQUE both of which admit $n^{1-\varepsilon}$ approximations in time $2^{\tilde{O}(n^\varepsilon)}$ (by partitioning the graph into sets of size n^ε and solving the problem optimally on these sets), despite the known $n^{1-o(1)}$ hardness of

³In fact, this will also show that the integrality gap remains at most $1 + O(\varepsilon)$ throughout the algorithm (as opposed to 2, which we used in the proof of Lemma 5.8), thus bounding the depth of the recursion by $O(1/\varepsilon^2)$ rather than $O(1/\varepsilon^3)$.

approximation for both problems [16, 13, 32]. This may also be taken as evidence that the recent subexponential time algorithm of Arora, Barak, and Steurer [3] for UNIQUE GAMES does not necessarily imply that UNIQUE GAMES is not hard, or not as hard as other NP-hard optimization problems.

Finally, turning to lift-and-project methods, we note here that our LS_+ -based algorithm for KNAPSACK shows that in some instances reduced integrality gaps which rely heavily on properties of the Lasserre hierarchy can be achieved using the weaker LS_+ hierarchy. This raises the question of whether the problems discussed in the recent series of Lasserre-based approximation algorithms [5, 15, 27] also admit similar results using LS_+ . On the flip side, it would also be interesting to see whether any such problems have strong integrality gap lower bounds for LS_+ , which would show a separation between the two hierarchies.

Acknowledgements We would like to thank Dana Moshkovitz for pointing out the blowup in her SET COVER reduction. We would also like to thank Mohammad R. Salavatipour for preliminary discussions on sub-exponential time approximation algorithms in general, and Claire Mathieu for insightful past discussions of the KNAPSACK-related results in [19]. Finally, we would like to thank Marek Cygan for bringing [11] to our attention.

References

- [1] M. Alekhnovich, S. Arora, and I. Turlakis. Towards strong nonapproximability results in the Lovász-Schrijver hierarchy. In *Proceedings of ACM Symposium on Theory of Computing*, pages 294–303, 2005.
- [2] N. Alon, D. Moshkovitz, and S. Safra. Algorithmic construction of sets for k -restrictions. *ACM Trans. Algorithms*, 2(2):153–177, April 2006.
- [3] S. Arora, B. Barak, and D. Steurer. Subexponential algorithms for unique games and related problems. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, pages 563–572, 2010.
- [4] S. Arora, B. Bollobás, L. Lovász, and I. Turlakis. Proving integrality gaps without knowing the linear program. *Theory of Computing*, 2(1):19–51, 2006.
- [5] B. Barak, P. Raghavendra, and D. Steurer. Rounding semidefinite programming hierarchies via global correlation. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, pages 472–481, 2011.
- [6] M. Charikar, K. Makarychev, and Y. Makarychev. Integrality gaps for Sherali-Adams relaxations. In *Proceedings ACM Symposium on Theory of Computing*, pages 283–292, 2009.
- [7] C. Chekuri and S. Khanna. A polynomial time approximation scheme for the multiple knapsack problem. *SIAM Journal on Computing*, 35(3):713–728, 2005.
- [8] E. Chlamtáč. Approximation algorithms using hierarchies of semidefinite programming relaxations. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, 0:691–701, 2007.

- [9] E. Chlamtáč and M. Tulsiani. Convex relaxations and integrality gaps. In Miguel F. Anjos and Jean B. Lasserre, editors, *Handbook on Semidefinite, Conic and Polynomial Optimization*, volume 166 of *International Series in Operations Research & Management Science*, pages 139–169. Springer US, 2012.
- [10] V. Chvátal. A greedy heuristic for the set-covering problem. *Mathematics of Operations Research*, 4(3):233–235, 1979.
- [11] M. Cygan, L. Kowalik, and M. Wykurz. Exponential-time approximation of weighted set cover. *Inf. Process. Lett.*, 109(16):957–961, July 2009.
- [12] U. Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, July 1998.
- [13] U. Feige and J. Kilian. Zero knowledge and the chromatic number. *Journal of Computer and System Sciences*, 57(2):187 – 199, 1998.
- [14] K. Georgiou, A. Magen, T. Pitassi, and I. Tourlakis. Integrality gaps of $2 - o(1)$ for vertex cover SDPs in the Lovász–Schrijver hierarchy. *SIAM Journal on Computing*, 39(8):3553–3570, 2010.
- [15] V. Guruswami and A.K. Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives. In *Proceedings of IEEE Symposium on Foundations of Computer Science*, pages 482 –491, oct. 2011.
- [16] J. Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Mathematica*, 182:105–142, 1999.
- [17] R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [18] D. S. Johnson. Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.*, 9(3):256–278, December 1974.
- [19] A. Karlin, C. Mathieu, and C. Nguyen. Integrality gaps of linear and semi-definite programming relaxations for knapsack. In *Proceedings of Conference on Integer Programming and Combinatorial Optimization*, pages 301–314, 2011.
- [20] R.M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103, 1972.
- [21] J. B. Lasserre. An explicit equivalent positive semidefinite program for nonlinear 0-1 programs. *SIAM Journal on Optimization*, 12(3):756–769, 2002.
- [22] L. Lovász. On the ratio of optimal integral and fractional covers. *Discrete Mathematics*, 13(4):383 – 390, 1975.
- [23] L. Lovász and A. Schrijver. Cones of matrices and set-functions and 0-1 optimization. *SIAM J. Optim.*, 1(2):166–190, 1991.
- [24] D. Moshkovitz. The projection games conjecture and the NP-hardness of $\ln n$ -approximating set-cover. In *Proceedings of APPROX*, pages 276–287, 2012.

- [25] D. Moshkovitz and R. Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, June 2008.
- [26] P. Raghavendra. Optimal algorithms and inapproximability results for every csp? In *Proceedings of ACM Symposium on Theory of Computing*, pages 245–254, 2008.
- [27] P. Raghavendra and N. Tan. Approximating csps with global cardinality constraints using sdp hierarchies. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, pages 373–387. SIAM, 2012.
- [28] H. D. Sherali and W. P. Adams. A hierarchy of relaxations between the continuous and convex hull representations for zero-one programming problems. *SIAM Journal on Discrete Mathematics*, 3(3):411–430, 1990.
- [29] M. Singh and K. Talwar. Improving integrality gaps via Chvátal-Gomory rounding. In *Proceedings of APPROX*, pages 366–379, 2010.
- [30] P. Slavík. A tight analysis of the greedy algorithm for set cover. *Journal of Algorithms*, 25(2):237 – 254, 1997.
- [31] A. Srinivasan. Improved approximations of packing and covering problems. In *Proceedings of ACM Symposium on Theory of computing*, pages 268–276, 1995.
- [32] D. Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007.

A A Combinatorial Approximation Algorithm

The following result is well-known and follows from standard dual-fitting techniques.

Theorem A.1 (Chvátal [10], paraphrased). *The greedy algorithm is a polynomial-time H_b -approximation where b is the size of the largest cover-set in \mathcal{S} . Moreover, it finds a solution whose cost is at most H_b times the optimum value of the standard LP relaxation.*

Motivated by the former Theorem, we present next our combinatorial algorithm (sketched in Section 3.1) in full details.

The outer loop of Algorithm 2 is reminiscent of the steps in many PTASes that “guess” the largest items to be used (for example, many KNAPSACK variants such as in [7]). This is almost a correct interpretation of the guesswork done in Algorithm 2. The following simple structural result will help articulate this notion in the proof of the main result.

Lemma A.2. *If \mathcal{C} is solution to a SET COVER instance (X, \mathcal{S}) , then there is an ordering S_1, \dots, S_k of the cover-sets in \mathcal{C} such that $|S_i \setminus \bigcup_{j=1}^{i'-1} S_j| \leq \frac{n}{i'}$ for any $1 \leq i' \leq i \leq k$.*

Proof. Order \mathcal{C} greedily: Iteratively select the set in \mathcal{C} that covers the most uncovered items. Then, by definition of this ordering, for the first i' sets we have $n \geq |\bigcup_{j=1}^{i'} S_j| \geq i' \cdot |S_{i'} \setminus \bigcup_{j=1}^{i'-1} S_j|$. Moreover, for $i > i'$, since $S_{i'}$ is chosen instead of S_i , we must have $|S_i \setminus \bigcup_{j=1}^{i'-1} S_j| \leq |S_{i'} \setminus \bigcup_{j=1}^{i'-1} S_j| \leq \frac{n}{i'}$. \square

Algorithm 2 An Improved Approximation for SET COVER

```
1:  $\mathcal{C} \leftarrow \mathcal{S}$ 
2: for each collection  $\mathcal{D}$  of at most  $d$  sets in  $\mathcal{S}$  do
3:   Let  $\mathcal{T} \leftarrow \{S \in \mathcal{S} : |S \setminus \bigcup_{T \in \mathcal{D}} T| \leq \frac{n}{d}\}$ 
4:   if  $\mathcal{T} \cup \mathcal{D}$  cannot cover  $X$  then skip to the next iteration
5:   while  $\mathcal{D}$  does not cover  $X$  do
6:     Let  $S \leftarrow \arg \min_{S \in \mathcal{T}} \frac{c(S)}{|S \setminus \bigcup_{T \in \mathcal{D}} T|}$ 
7:      $\mathcal{D} \leftarrow \mathcal{D} \cup \{S\}$ 
8:   end while
9:   if  $\mathcal{D}$  is cheaper than  $\mathcal{C}$  then  $\mathcal{C} \leftarrow \mathcal{D}$ 
10: end for
11: return  $\mathcal{C}$ 
```

We are now ready to give all details for the combinatorial proof of Theorem 1.1.

Proof of Theorem 1.1. Let \mathcal{C}^* be an optimal solution of cost OPT and size, say, k . Furthermore, let S_1^*, \dots, S_k^* be an ordering of the cover-sets in \mathcal{C}^* with the property guaranteed by Lemma A.2. If $k \leq d$ then the set \mathcal{C}^* will be considered in some iteration so Algorithm 2 will actually find an optimal solution.

Otherwise, consider the iteration of the outer loop that guessed $\mathcal{D} = \{S_1^*, \dots, S_d^*\}$. Consider the SET COVER instance (Y, \mathcal{T}) with items $Y := X \setminus \bigcup_{T \in \mathcal{D}} T$ and cover-sets $\mathcal{T} := \{S \in \mathcal{S} : |S \setminus \bigcup_{T \in \mathcal{D}} T| \leq \frac{n}{d}\}$. The cost of the cover-sets $S \in \mathcal{T}$ in this instance should be equal to their original costs (before we subtracted $\bigcup_{T \in \mathcal{D}} T$). Each cover-set in \mathcal{T} has size at most $\frac{n}{d}$ so, by Theorem A.1, the greedy algorithm applied to this instance is an $H_{n/d}$ -approximation. Lemma A.2 guarantees that the restriction of the cover-sets S_{d+1}^*, \dots, S_k^* to $X \setminus \bigcup_{T \in \mathcal{D}} T$ have size at most $\frac{n}{d}$ so these restrictions form a valid solution for (Y, \mathcal{T}) of cost $\sum_{i=d+1}^k c(S_i^*)$. Thus, the optimum solution to the SET COVER instance (Y, \mathcal{T}) has cost $\sum_{i=d+1}^k c(S_i^*)$. Steps 5-8 act just like the greedy algorithm on this instance (Y, \mathcal{T}) which, by Theorem A.1, adds a total of $H_{\frac{n}{d}} \cdot \sum_{i=d+1}^k c(S_i^*)$ to the cost of \mathcal{D} . The overall cost of the final solution \mathcal{D} is at most $\sum_{i=1}^d c(S_i^*) + H_{\frac{n}{d}} \cdot \sum_{i=d+1}^k c(S_i^*) \leq H_{\frac{n}{d}} \cdot \sum_{i=1}^k c(S_i^*) = H_{\frac{n}{d}} \cdot \text{OPT}$. Finally, note that there are $m^{O(d)}$ iterations of the outer loop, each taking polynomial time. \square